

Evaluation of Using Inductive/Capacitive-Coupling Vertical Interconnects in 3D Network-on-Chip

Jin Ouyang*, Jing Xie, Matthew Poremba, Yuan Xie

Department of Computer Science and Engineering, the Pennsylvania State University

*jouyang@cse.psu.edu

Abstract—In recent 3DIC studies, through silicon vias (TSV) are usually employed as the vertical interconnects in the 3D stack. Despite its benefit of short latency and low power, forming TSVs adds additional complexities to the fabrication process. Recently, inductive/capacitive-coupling links are proposed to replace TSVs in 3D stacking because the fabrication complexities of them are lower. Although state-of-the-art inductive/capacitive-coupling links show comparable bandwidth and power as TSV, the relatively large footprints of those links compromise their area efficiencies. In this work, we study the design of 3D network-on-chip (NoC) using inductive/capacitive-coupling links. We propose three techniques to mitigate the area overhead introduced by using these links: (a) *serialization*, (b) *in-transceiver data compression*, and (c) *high-speed asynchronous transmission*. With the combination of these three techniques, evaluation results show that the overheads of all aspects caused by using inductive/capacitive-coupling vertical links can be bounded under 10%.

I. INTRODUCTION

In recent years a shift of researches from merely technology scaling to massive integration has been witnessed in the semiconductor industry. Three-dimensional integrated circuits (3DICs) bring brand new opportunities to integrate many and heterogeneous components in a single chip, to deliver performance and functionalities hard to rival by 2D chips and offer unique opportunities to continue the Moore's law [11].

Despite the merits mentioned above, the additional complexities introduced by 3DIC to the fabrication process cast a shadow on the picture. In TSV based 3DIC, forming vertical links to interconnect different layers requires additional and unconventional processing steps involving high aspect-ratio etching, wafer thinning, bonding, etc. Defects that occur during the extra processes cause performance and yield loss, hindering the wide adoption of 3DIC.

This hurdle of fabrication has led to the advent of AC-coupling vertical links recently. For this kind of links, either magnetic flux (inductive-coupling [5], [7], [8], [12]) or electric field (capacitive-coupling [2]–[4]) serve as the medium through which signals are transmitted vertically. The advantage of inductive/capacitive coupling links is that the transceivers can be implemented by pure 2D process. Additional 3D processing steps to form TSVs are eliminated. In addition, it is possible to improve the performance of inductive/capacitive-coupling links with technology scaling. However, Most reported inductive/capacitive-coupling links

This work is supported in part by NSF 0905365, 0903432, and 0702617, and SRC grants.




Fig. 1. A conceptual view of vertical interconnects based on three different physical media

have much larger area than TSV, ranging from hundreds to tens of thousands of square microns. The inferior area efficiency of inductive/capacitive-coupling links renders their applicability in high density 3DIC a question.

The focus of this paper is to improve the area efficiency of inductive/capacitive-coupling links, and evaluate their applicability in the context of 3D NoCs. To achieve this goal, three techniques are proposed:

- **Serial communication** on the vertical links is firstly proposed to constrain the area overheads. However, serial communication incurs performance penalty since the effective vertical bandwidth is reduced.
- **Data compression** is then proposed to recover the performance by reducing the bandwidth requirements for vertical links. Different from previous work, we implement an *in-transceiver data compression* scheme where compression is performed simultaneously with serial communication, incurring no latency overhead nor changes to original NoC.
- **High-speed asynchronous transmission** is lastly exploited to further boost the performance. Inductive/capacitive-coupling vertical links combined with serial transceivers become a natural asynchronous interface capable of running at a frequency much higher than the 2D links. This provides opportunities to further improve the area efficiency of vertical links.

A comprehensive evaluation of a 3D CMP interconnected by NoC using inductive/capacitive-coupling links is presented, taking into consideration performance, energy, and area. With the combination of serialization, data compression, and asynchronous transmission, experiment results show that the overheads of performance, energy, and area caused by using inductive/capacitive-coupling links can all be bounded under 10%. In addition, the area efficiency of these links is also brought to over 90% of TSV. Considering the benefit of reducing fabrication difficulties, this is a positive result to support inductive/capacitive-coupling communication as a competitive alternative in real 3D designs. To the best of our knowledge, this is the first work to evaluate the utility

TABLE I. Recent Work on Inductive/Capacitive-Coupling Vertical Links

	Inductive					Capacitive				
	Xu05 [12]	Ishikuro07 [5]	Miura08 [7]	Miura09 [8]		Sun07 [4]	Fazzi07 [2]	Gu07 [3]		
Energy (pJ/b)	17	>390	0.14	1.4	1.5	1.8	3	0.08	0.27	0.39
Data Rate (GB/s)	2.8	0.02	1	11		1.8	1.7-2.46	10	11	
Inductor/Capacitor Size	150×150um ²	>0.6mm ²	29um (diameter)	120um (diameter)		NA	8×8-20×20um ²	48×18um ²	72×38um ²	
Inductor/Capacitor Pitch	200um	NA	30um	NA		NA	NA	NA		
Active Area (mm ²)	NA	>0.006	NA	NA		NA	NA	0.0012	0.0021	
Total Area (mm ²)	NA	NA	NA	0.0015		NA	NA	NA		
Distance (um)	90	1200	15	15	30	45	<10	1	3	
Process	.35um	.25um	90nm	180nm		180nm	130nm	180nm		

TABLE II. [11]

TSV Characteristics	
Size	10×10um ²
Height	50um
Pitch	20um
R _{TSV}	14mΩ
C _{TSV}	284fF
Delay	60.6ps
Energy	0.125pJ

of inductive/capacitive interconnects in 3D NoC and propose to combine all the three techniques above to improve their performance.

II. VERTICAL INTERCONNECTIONS IN 3DIC

A. Through Silicon Vias

TSVs are the most commonly used vertical interconnects. Usually the dimensions of TSV are quite small, resulting in low parasitics which in turn translate into high bandwidth and low energy consumption. The specifications of TSV depend on the process and can vary diversely. For example, IBM, IMEC, and many more organizations have fabricated TSVs with different pitches ranging from several microns up to tens of microns [11]. In this paper we assume a moderate TSV whose dimensions and RC data is generalized in Table II [11]. The delay and energy results are obtained from Spice simulation based on the methodologies in by using [1].

B. AC-Coupling Vertical Links

Table I summarizes most recently demonstrated designs of both inductive-coupling links (column 2-5) and capacitive-coupling links (column 6-8).

Due to the different natures of the carrier fields, the characteristics of inductive-coupling and capacitive-coupling links also differ. The communication distance of inductive-coupling is relatively long. This long communication distance is exploited in [5] to realize an attachable wireless chip probe and most recently in [9] for a memory-on-processor chip where the communication distance is 120um. Capacitive-coupling, on the other hand, has much shorter communication distance which constrains it to be used only in face-to-face stacking. However, the footprints of capacitive-coupling are likely to be much smaller than inductive-coupling, incurring much lower area overhead than inductive-coupling.

Existing inductive-coupling and capacitive-coupling designs can achieve a data rate as high as 11Gb/s or energy consumption as low as 0.14pJ/b, which are both comparable to TSV (the first and second rows in Table I). In addition, they can be

TABLE III. Implementation Results of 128-bit Transceiver with STMicro 65nm technology at 5GHz

Transceiver w/o Data Compression		
Serialization Ratio	Energy (pJ/b)	Area (um ²)
4:1	0.3581	4776
8:1	0.5987	4102
16:1	0.9063	3890
Transceiver w Data Compression		
Serialization Ratio	Energy (pJ/b)	Area (um ²)
4:1	0.7785	8473
8:1	0.9528	7009
16:1	1.4801	6334

implemented with conventional 2D process without additional processing steps. The metal solenoids and capacitive plates can be formed simply by the top metal of the integrated circuit. However, as Table I shows (the fifth and sixth rows), most reported inductive/capacitive-coupling links have much larger area than TSV, ranging from hundreds of square microns to tens of thousands of square microns.

In the remaining evaluations of the paper, the characteristics of inductive/capacitive-coupling links are based on Table I where more conservative processes (90nm or older) are used. However, for the transceivers, and the 3D NoC and the CMP introduced in the following sections, a more contemporary technology node of 65nm is assumed. Hence, our evaluations hereafter can be deemed as the *worst-case* results, since it is possible to improve AC-coupling links' performance with technology scaling [8]. Moreover, due to the incompleteness and large variation of reported data on inductive/capacitive-coupling links, we design our evaluations to consider a range of possible metrics of AC-coupling links.

III. REDUCING AREA OVERHEAD WITH SERIALIZATION

One way to address the large area overhead of inductive/capacitive-coupling vertical links is to adopt serial communication where the data is transmitted on a single link. The *serialization ratio* is defined as the ratio between the number of parallel and serial links.

We adopt the transceiver design proposed in [10] originally for TSV, shown in Figure 2a and 2b. This transceiver effectively implements the *start-1 stop-0 serial protocol*. Note that the transceiver owns private clock generators (ring-oscillator) which allows for asynchronous transmission. We implement the transceiver with STMicroelectronics 65nm technology. The transceiver can run at as high as 5Ghz using this technology. The first half of Table III shows power and area of 128-bit

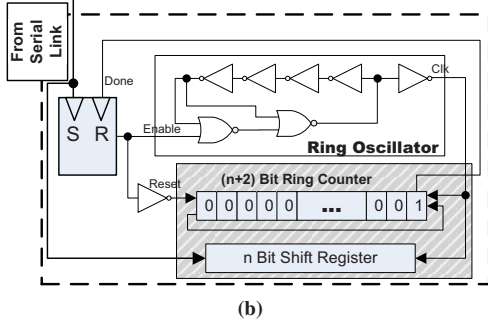
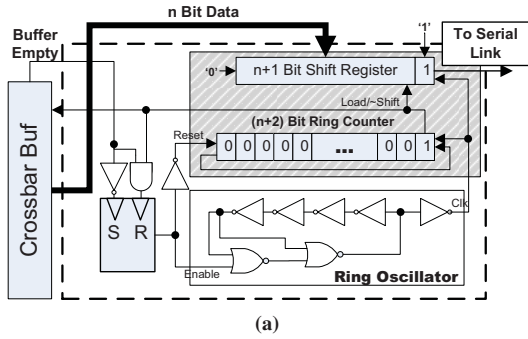


Fig. 2. (a) A 7-port 3D NoC router with dedicated vertical ports (Up/Down). (b) Serial transmitter. (c) Serial receiver.

transceivers (transmitter + receiver) with different serialization ratios (due to space limits, Table III only lists results for 4:1, 8:1, and 16:1 transceivers). Note that a 128-bit transceiver may consist of multiple narrower *sub*-transceivers, say, 32 4-bit *sub*-transceivers.

The impact of serialization on the area of an NoC router is evaluated based on the implementation results. We derive the area and energy of the 3D NoC router with ORION 2.0 [6]. Figure 3 shows the total router area including area taken by vertical links, normalized to the router area without vertical links. As a guidance, Figure 3 also plots relative bandwidth reduction with serialization. We notice that with the large pitches of inductive/capacitive-coupling links, up to two times of area overhead is observed without serialization. In addition, we see that with a serialization ratio of 8:1, the area of the router with inductive/capacitive links is only 1.1 times larger than the router with TSV for pitches not larger than 100 μ m.

IV. DATA COMPRESSION

As shown by Fig. 3 and subsequently in Section V, serialization can incur considerable performance penalty. In this section, we propose to use data compression to reduce the performance penalty and develop an *in-transceiver data compression* design.

A. The Concept

In this work we carry out zero-pattern compression at the granularity of single 32-bit words. Conceptually, if one word in a flit contains all zeros, this word is omitted for transmission and instead only a few bits to encode this situation are sent. For instance, consider a 128-bit wide flit that contains four




Fig. 3. Area overheads when the serialization ratio varies, for different vertical link pitches. The results have been normalized to the case where the area of vertical links is not included. The trend line shows the normalized bandwidth reduction with increasing serialization ratio.




Fig. 4. Zero-pattern compression in principle.

32-bit words (Fig. 4). If some word in the flit contains all zeros, it is eliminated and only the remaining three words are sent. Encoding can be done in hardware by generating a 4-bit mask. A “1” in the mask represents a valid word while a “0” indicates the corresponding word is omitted. This mask is sent together with the compressed data, and used by the receiver to recover the data.

B. Implementation

In this subsection we use an 128-bit 4:1 transceiver to illustrate the implementation of data-compression enabled transceiver. Except for some minor changes, the implementation is similar for other serialization ratios. The synthesis results of the enhanced transceivers with 65nm technology are listed in the second half of Table III.

1) *Data Compression Operations*: Fig. 5 shows the additional operations needed by data compression during each transmission cycle. In Fig. 5a, simultaneously with loading data into shift register, the transmitter also generates and latches the 4-bit mask into a register. This mask is sent on parallel links, together with the start-bit on the serial link in the second cycle, and each valid data bit is sent one-by-one in following cycles. At the end of the transaction, all valid data bits are sent, and the mask register is reset. In Fig. 5b, when detecting a rising-edge on the serial link, the receiver latches the start-bit into its data register and the mask on the parallel links into its mask register. Then when following bits come in, the mask register is used to rearrange the bits to recover the original data. At the end of the transaction, the decompressed data is delivered and the mask register is cleared.

2) *Hardware Details*: Supporting data compression only adds a few changes to the transceiver shown in Fig. 2.




Fig. 5. Operations in a single transaction. The text in shaded boxes indicates additional operations of data compression.

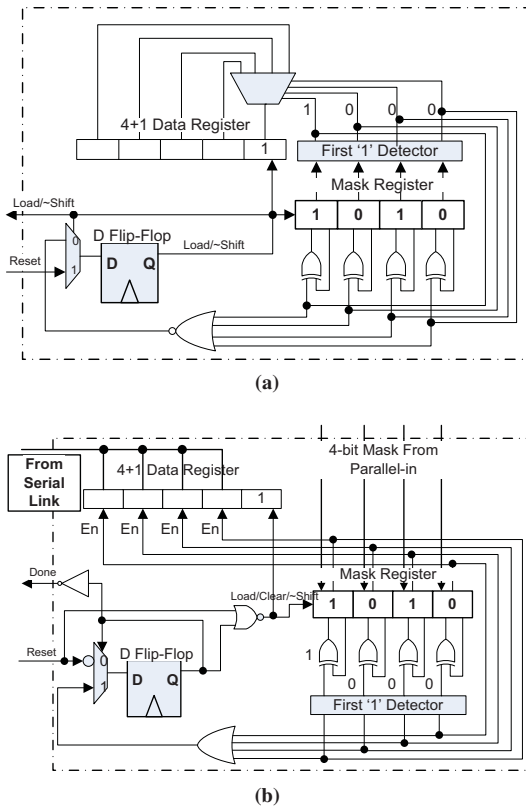


Fig. 6. Data compression hardware details: (a) Compression. (b) Decompression.

Specifically, only the shaded parts in Fig. 2a and 2b are modified.

The modified shaded part in Fig. 2a is shown in Fig. 6a. For clarity, clock and data signals are omitted. As can be seen, there are two major changes:

- The ring-counter is replaced by the mask register, combined with some logic, to control the transmission sequence.
- The data register is not a pure shift register, but instead is equipped with a 4-to-1 multiplexer to forward any bits to the head of the register.

The data and mask registers are loaded when crossbar buffer becomes ready, as shown. In each following transmission cycle, with a first-one detector, the mask register controls the multiplexer to select the next valid bit to be sent. The output from the first-one detector is also fed-back to clear the leading “1” in the mask register in each cycle. After all valid bits are sent, the mask register should contain only “0”s and reset the transmitter state. The modifications to the receiver, as shown in Fig. 6b, are very similar to the transmitter. The difference is that instead of using a multiplexer, the data register demultiplexes the incoming serial data using data enables generated by the mask register. One subtlety with receiver is that the data register needs to be cleared to ‘0’ before receiving the first valid bit, to eliminate obsolete bits from last transaction.

V. EXPERIMENT

In this section, inductive/capacitive-coupling vertical links are evaluated in the context of a 3D NoC with a cycle accurate simulator. In the simulator we model a 2-tier, 16-core 3D CMP with a shared L2 cache distributed into 16 banks. These 32 nodes are interconnected by a 3D NoC configured as a 3D MESH (4×4 by 2 layers). We choose a mosaic floorplan where each core is surrounded by 5 L2 cache banks in the cardinal directions (except for cores on edges). The NoC router metrics are derived with ORION 2.0 [6]. The area and energy of proposed transceivers is given in Table III, for STMicro 65nm technology and 5GHz frequency. The whole CMP (processor, cache, NoC) is clocked at 1.5Ghz. For synthetic traffics, we study the impact of serialization and asynchronous transmission; for multi-thread applications, we evaluate all the three proposed approaches.

Note: For comparison purpose, we define $baseline_{TSV}$ as the case that TSV is used without serialization. As discussed in Section III, serialization with TSV only gains negligible area savings, but incurs significant performance loss. The two proposed approaches (data compression and asynchronous transmission) to recover performance are only effective when serialization is used, where vertical link bandwidth is the bottleneck. **Therefore, none of serialization, data compression, and asynchronous transmission is used in $baseline_{TSV}$.**

A. Study with Synthetic Traffics

Fig. 7 shows the average message latencies for four synthetic traffics with different load ratios and serialization ratios. Note that the performance when the serialization ratio is 1:1 is just the same as $baseline_{TSV}$. It can be seen from the results using serialization with the vertical links negatively affects the performance of NoC. With the serialization ratio increasing, the message latency under the same load ratio increases, and the saturation load ratio decreases. Within the four traffic

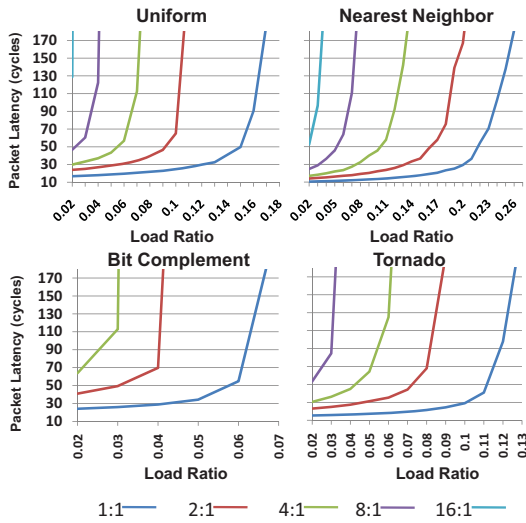


Fig. 7. Message latencies for different traffic patterns without data compression.

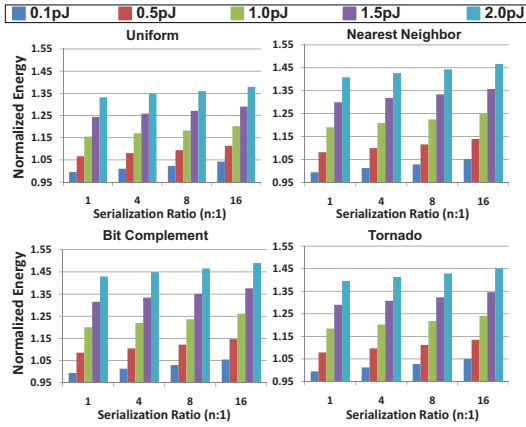


Fig. 8. Energy consumption normalized to $baseline_{TSV}$ for different traffic patterns.

patterns, *nearest neighbor* shows least performance reduction, while *bit complement* shows worst performance loss.

The energy consumption of the interconnection network normalized to $baseline_{TSV}$ is shown in Fig. 8. Since the energy consumption of inductive/capacitive-coupling vertical links can range from 0.1pJ to several hundred pJ, we study how the energy overhead changes over a range of link energies. This figure shows that both high serialization ratio and large inductive/capacitive-coupling link energy increase the energy overhead as compared to TSV. In addition, inductive/capacitive-coupling link energy affects the energy




Fig. 9. Message latency reduction with higher transmission frequencies.




Fig. 10. Effectiveness of data compression in improving IPC.




Fig. 11. Impacts of higher transmission frequency on IPC.

overhead to a greater extent. In order to constrain the overhead under 10% at 8:1 serialization ratio, the link energy has to be less than 0.5pJ/b.

Then, we study the effectiveness of high-speed asynchronous transmission in recovering performance loss. In this evaluation, a moderate load ratio of 0.05, and an serialization ratio of 4:1 are assumed. The message latencies with different transmission frequencies are plotted in Fig. 9 and normalized to the $baseline_{TSV}$. The y -axis of this figure only ranges from 1 to 3, as we consider that latencies that are three times higher are already beyond saturation. It can be seen that with the successively increasing transmission frequency, the packet latency is continuously decreasing. At 5Ghz, the average latency overheads is reduced to 25%.

B. Study with Multi-Thread Applications

In this subsection, four multi-thread benchmarks (*sap*, *sjas*, *sjbb*, and *tpc*) are simulated. Each of the benchmarks runs 16 threads on the 16 cores. In every run we first warm up the cache states by 3 million instructions for each thread, and then collect statistics for the next 5 million instructions. The IPCs of the benchmarks both without data compression

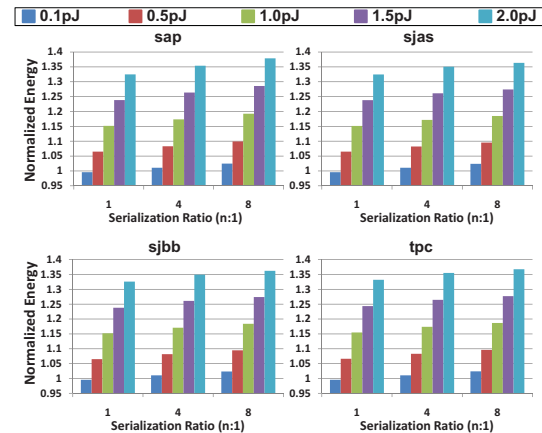


Fig. 12. Energy of interconnection network for multi-thread applications without data compression.

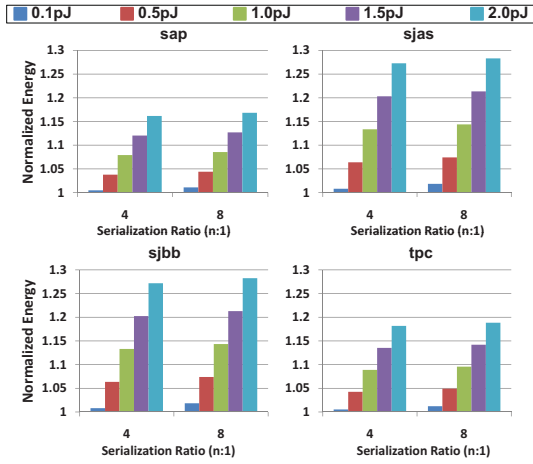


Fig. 13. Energy of interconnection network for multi-thread applications with data compression.




Fig. 14. Area efficiency of the interconnection network when different techniques are successively applied.

and with data compression are plotted in Fig. 10 and also normalized to the baseline. It can be seen that serialization does negatively impact the performance, and for a serialization ratio of 8:1, the performance is reduced to around 60% of original performance. With data compression used, this performance gap shrinks, especially for *sap* and *tpc* that have relatively high zero-pattern ratios.

The energy consumption of interconnection network normalized to $baseline_{TSV}$ is plotted in Fig. 12 without data compression, and in Fig. 13 with data compression. Without data compression, the energy overheads of the benchmarks are similar to that of synthetic traffics. With 0.5pJ/b link energy, the overhead can be limited under 10%. With data compression, on the other hand, the average energy overhead is reduced to 5% with 0.5pJ/b link energy.

Although effective, data compression is not able to recover the performance to satisfaction (only 70% to 75% in the 8:1 case). Asynchronous transmission is then exploited to further reduce the performance gap. For the serialization ratio of 8:1, we improve the transmission frequency to 2, 3, 4, and 5GHz respectively. The results normalized to the baseline are shown in Fig. 11, which validate that increasing transceiver frequency improves the IPC. At 5GHz, the performance losses of the benchmarks have been reduced to under 6%.

C. Area Efficiency

For a more comprehensive evaluation, we study the area efficiencies of the interconnection network in terms of the performance of multi-thread applications. We define the area

efficiency to be $(IPC/Router_Area)$. Three cases are compared: (a) *no data compression nor asynchronous transmission* is used; (b) *only data compression* is used; and (c) *both data compression and asynchronous transmission* are used. For all cases, the serialization ratio is 8:1. For asynchronous transmission, the transmission frequency is set to 5GHz. The inductive/capacitive-coupling link pitch is set to 70um, a moderate value among existing designs. The area efficiencies of using inductive/capacitive-coupling vertical links are normalized to $baseline_{TSV}$.

The results are shown in Fig. 14. Without data compression and asynchronous transmission, the average area efficiency of the four benchmarks is only 55% of $baseline_{TSV}$. When data compression is applied, the average area efficiency is increased to 67%. Finally, the area efficiencies of all four benchmarks are improved to over 90% when both data compression and asynchronous transmission are used, with an average of 92%.

VI. CONCLUSION

In this work we propose design methodologies for using inductive/capacitive-coupling vertical links in 3D NoC. To address the area overheads of these links, *serial communication* is adopted and it is shown that to reduce the overhead under 10%, the serialization ratio has to be at least 8:1. We design and implement asynchronous transceivers that can do in-place zero-pattern compression simultaneously with serialization. As a result of the three techniques, the performance of multi-thread applications are improved to over 94% of the baseline case, and the area efficiency is brought to above 90% of TSV. Finally, the energy overheads can be bounded under 10% with a link energy less than 0.5pJ/b

REFERENCES

- [1] S. M. Alam et al. Inter-strata connection characteristics and signal transmission in three-dimensional (3D) integration technology. In *Proc. of ISQED '07*, pages 580–585, 2007.
- [2] A. Fazzi et al. 3D capacitive interconnections with mono- and bi-directional capabilities. In *Proc. of ISSCC '07*, pages 356–608, Feb. 2007.
- [3] Q. Gu et al. Two 10Gb/s/pin low-Power interconnect methods for 3D ICs. In *Proc. of ISSCC '07*, pages 448–614, Feb. 2007.
- [4] D. Hopkins et al. Circuit techniques to enable 430Gb/s/mm² proximity communication. In *Proc. of ISSCC '07*, pages 368–609, Feb. 2007.
- [5] H. Ishikuro et al. An attachable wireless chip access interface for arbitrary data rate using pulse-based Inductive-coupling through LSI package. In *Proc. of ISSCC '07*, pages 360–608, Feb. 2007.
- [6] A. B. Kahng et al. ORION2.0: A fast and accurate noc power and area model for early-stage design space exploration. Technical report, Sep. 2008.
- [7] N. Miura et al. A 0.14 pJ/b inductive-coupling transceiver with digitally-controlled precise pulse shaping. *Journal of Solid-State Circuits*, 43(1):285–291, Jan. 2008.
- [8] N. Miura et al. A high-speed inductive-coupling link with burst transmission. *Journal of Solid-State Circuits*, 44(3):947–955, March 2009.
- [9] K. Niitsu et al. An inductive-coupling link for 3D integration of a 90nm CMOS processor and a 65nm CMOS SRAM. In *Proc. of ISSCC '09*, pages 480–481, Feb. 2009.
- [10] S. Pasricha. Exploring serial vertical interconnects for 3D ICs. In *Proc. of DAC '09*, pages 581–586, Jul. 2009.
- [11] Y. Xie et al. *Three-Dimensional Integrated Circuit Design*. Springer, 2010.
- [12] J. Xu et al. 2.8 Gb/s inductively coupled interconnect for 3D ICs. In *Proc. of ISVLSI '05*, pages 352–355, Jun. 2005.